arXiv:1405.0338v1 [math.ST] 2 May 2014

Rate Optimal Denoising of Simultaneously Sparse and Low Rank Matrices *

Dan Yang[1] Zongming Ma[2] and Andreas Buja[2]

[1]Rutgers University and [2]University of Pennsylvania

Abstract

We study minimax rates for denoising simultaneously sparse and low rank matrices in high dimensions. We show that an iterative thresholding algorithm achieves (near) optimal rates adaptively under mild conditions for a large class of loss functions. Numerical experiments on synthetic datasets also demonstrate the competitive performance of the proposed method.

Keywords: Denoising, High dimensionality, Low rank matrices, Minimax rates, Simultaneously structured matrices, Sparse SVD, Sparsity.

1 Introduction

In recent years, there has been a surge of interest in estimating and denoising structured large matrices. Leading examples include denoising low rank matrices [10], recovering low rank matrices from a small number of entries, i.e., matrix completion [7, 6, 13, 14, 20], reduced rank regression [3], group sparse regression [29, 17], among others.

In the present paper, we study the problem of denoising an $m \times n$ data matrix

$$\mathbf{X} = \mathbf{M} + \mathbf{Z}. \tag{1}$$

The primary interest lies in the matrix $\mathbf{M}$ that is sparse in the sense that nonzero entries are assumed to be confined on a $k \times l$ block, which is not necessarily consecutive. In addition to being sparse, the rank of $\mathbf{M}$, denoted by r, is assumed to be low. Thus, $\mathbf{M}$ can be regarded as *simultaneously structured* as opposed to those simply structured cases where $\mathbf{M}$ is assumed to be either only sparse or only of low rank. To be concrete, we assume that $\mathbf{Z}$ consists of i.i.d. additive Gaussian white noise with variance σ^2. In the literature, the problem has also been referred to as the sparse SVD (singular value decomposition) problem. See, for instance, [28] and the references therein.

The interest in this problem is motivated by a number of related problems:

1. *Biclustering.* It provides an ideal model for studying biclustering of microarray data. Let the rows of $\mathbf{X}$ correspond to cancer patients and the columns correspond to gene expression levels measured with microarrays. A subset of k patients can be clustered together as a subtype of the same cancer, which in turn is determined by a subset of l genes. Moreover, the gene expression levels on such a bicluster can usually be captured by a low rank matrix. See, e.g., Shabalin et al. [23], Lee et al. [15], Butucea and Ingster [5], Sun and Nobel [26], Chen et al. [8].

*Some partial results of this paper have been presented at the 2013 Allerton Conference as an invited paper [2].

2. *Recovery of simultaneously structured matrices with compressive measurements.* There has been emerging interest in the signal processing community in recovering such simultaneously structured matrices based on compressive measurements, partly motivated by problems such as sparse vector recovery from quadratic measurements and sparse phase retrieval. See, e.g., [24, 16] and the references therein. The connection between the recovery problem and the denoising problem considered here is partially explored in [21]. An interesting phenomenon in the recovery setting is that convex relaxation approach no longer works well [22] as it does in the simply structured cases.

3. *Sparse reduced rank regression.* The denoising problem is also closely connected to prediction in reduced rank regression where the coefficient matrix is also sparse. Indeed, let $n = l$, then problem (1) reduces to sparse reduced rank regression with orthogonal design. See [4] and [19] for more discussion.

The main contribution of the present paper includes the following: i) We provide information-theoretic lower bounds for the estimation error of $\mathbf{M}$ under squared Schatten-q norm losses for all $q \in [1, 2]$; ii) We propose a computationally efficient estimator that, under mild conditions, attains high probability upper bounds that match the minimax lower bounds within a multiplicative log factor (and sometimes even within a constant factor) simultaneously for all $q \in [1, 2]$. The theoretical results are further validated and supported by numerical experiments on synthetic data.

The rest of the paper is organized as follows. In Section 2, we precisely formulate the denoising problem and propose a denoising algorithm based on the idea of iterative thresholding. Section 3 establishes minimax risk lower bounds and high probability upper bounds that match the lower bounds within a multiplicative log factor for all squared Schatten-q norm losses with $q \in [1, 2]$. Section 4 presents several numerical experiments which demonstrate the competitive finite sample performance of the proposed denoising algorithm. The proofs of the main results are presented in Section 5, with some technical details relegated to Appendix A.

2 Problem Formulation and Denoising Method

Notation For any $a, b \in \mathbb{R}$, let $a \wedge b = \min(a, b)$ and $a \vee b = \max(a, b)$. For any two sequences of positive numbers $\{a_n\}$ and $\{b_n\}$, we write $a_n = O(b_n)$ if $a_n \le Cb_n$ for some absolute positive constant C and all n. For any matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, denote its successive singular values by $\sigma_1(\mathbf{A}) \ge \cdots \ge \sigma_{m\wedge n}(\mathbf{A}) \ge 0$. For any $q \in [1, \infty)$, the Schatten-q norm of $\mathbf{A}$ is defined as $\|\mathbf{A}\|_{\mathrm{S}_q} = (\sum_{i=1}^{m\wedge n} \sigma_i^q(\mathbf{A}))^{1/q}$. Thus, $\|\mathbf{A}\|_{\mathrm{S}_1}$ is the nuclear norm of $\mathbf{A}$ and $\|\mathbf{A}\|_{\mathrm{S}_2} = \|\mathbf{A}\|_{\mathrm{F}}$ is the Frobenius norm. In addition, the Schatten-∞ norm of $\mathbf{A}$ is $\|\mathbf{A}\|_{\mathrm{S}_\infty} = \sigma_1(\mathbf{A}) = \|\mathbf{A}\|_{\mathrm{op}}$, where $\|\cdot\|_{\mathrm{op}}$ stands for the operator norm. The rank of $\mathbf{A}$ is denoted by $\mathsf{rank}(\mathbf{A})$. For any vector $\mathbf{a}$, we denote its Euclidean norm by $\|\mathbf{a}\|$. For any integer m, $[m]$ stands for the set $\{1, \ldots, m\}$. For any subset $I \subset [m]$ and $J \subset [n]$, we use $\mathbf{A}_{IJ}$ to denote the submatrix of $\mathbf{A}$ with rows indexed by I and columns by J. When either I or J is the whole set, we replace it with $*$. For instance, $\mathbf{A}_{I*} = \mathbf{A}_{I[n]}$. Moreover, we use $\mathrm{supp}(\mathbf{A})$ to denote the set of nonzero rows of $\mathbf{A}$. For any set A, $|A|$ denotes its cardinality and A^c denotes its complement. A matrix $\mathbf{A}$ is called orthonormal, if the column vectors are of unit length and mutually orthogonal. For any event E, we use $\mathbf{1}_E$ to denote the indicator function on E, and E^c denotes its complement.

2.1 Problem Formulation

We now put the denoising problem in a decision-theoretic framework. Recall model (1). We are interested in estimating $\mathbf{M}$ based on the noisy observation $\mathbf{X}$, where $\mathbf{M}$ is simultaneously

sparse and low rank. Let the singular value decomposition (SVD) of $\mathbf{M} = \mathbf{UDV}'$, where $\mathbf{U}$ is $m \times r$ orthonormal, $\mathbf{V}$ is $n \times r$ orthonormal and $\mathbf{D} = \mathrm{diag}(d_1, \ldots, d_r)$ is $r \times r$ diagonal with $d_1 \geq \cdots \geq d_r > 0$. In addition, since the nonzero entries on $\mathbf{M}$ concentrate on a $k \times l$ block, $\mathbf{U}$ has at most k nonzero rows and $\mathbf{V}$ at most l. Therefore, the parameter space of interest can be written as

$$\begin{aligned} \mathcal{F}(m, n, k, l, r, d, \kappa) = \{\mathbf{M} = \mathbf{UDV}' \in \mathbb{R}^{m \times n} : \mathsf{rank}(\mathbf{M}) = r, \\ |\mathrm{supp}(\mathbf{U})| \leq k, |\mathrm{supp}(\mathbf{V})| \leq l, \\ d \leq d_r \leq \cdots \leq d_1 \leq \kappa d\}. \end{aligned} \tag{2}$$

We will focus on understanding the dependence of the minimax estimation error on the key model parameters (m, n, k, l, r, d), while $\kappa > 1$ is treated as an unknown universal constant. Without loss of generality, we assume $m \geq n$ here and after. Note that it is implicitly assumed in (2) that $m \geq k \geq r$ and $n \geq l \geq r$.

To measure the estimation accuracy, we use the following squared Schatten-q norm loss functions:

$$L_q(\mathbf{M}, \widehat{\mathbf{M}}) = \|\widehat{\mathbf{M}} - \mathbf{M}\|_{\mathrm{s}_q}^2, \qquad q \in [1, 2]. \tag{3}$$

The model (1), the parameter space (2) and the loss functions (3) give a precise formulation of the denoising problem.

2.2 Approach

From a matrix computation viewpoint, if one seeks a rank r approximation to a matrix $\mathbf{X}$, then one can first find its left and the right r leading singular vectors. If we organize these vectors as columns of the left and the right singular vector matrices $\mathbf{U}$ and $\mathbf{V}$, then the matrix $(\mathbf{UU}')\mathbf{X}(\mathbf{VV}')$ has the minimum Frobenius reconstruction error for $\mathbf{X}$ among all rank r matrices, since $\mathbf{U}'\mathbf{XV}$ will be a diagonal matrix consisting of the r leading singular values of $\mathbf{X}$. On the other hand, if one wants to enforce sparsity in the resulting matrix, it is natural to utilize the idea of thresholding in the above calculation. Motivated by the above observation and also by an iterative thresholding idea previously used in solving sparse PCA problem [18, 30], we propose the denoising scheme in Algorithm 1 via two-way iterative thresholding.

Without the two thresholding steps, the iterative part of the algorithm computes the leading singular vectors of any rectangular matrix, and can be viewed as a two-way generalization of the power iteration [12].

In the thresholding steps, we apply row-wise thresholding to the matrix $\mathbf{U}^{(t),\mathrm{mul}}$ (resp. $\mathbf{V}^{(t),\mathrm{mul}}$) obtained after the multiplication step. In the thresholding function $\eta(x, t)$, the second argument $t > 0$ is called the threshold level. In Algorithm 1, the first argument x is always non-negative. In order for the later theoretical results to work, we impose the following minimal assumption on the thresholding function η:

$$\begin{aligned} |\eta(x, t) - x| \leq t, \quad \text{for any } x \geq 0,\ t > 0, \\ \eta(x, t) = 0, \quad \text{for any } t > 0,\ x \in [0, t]. \end{aligned} \tag{4}$$

Examples of such thresholding functions include the usual soft and hard thresholding, the SCAD [11], the MCP [31], etc. Thus, for instance, when thresholding $\mathbf{U}^{(t),\mathrm{mul}}$, if η is the hard thresholding function, then we are going to keep all the rows whose norms are greater than γ_u and kill all the rows whose norms are smaller than γ_u. For other thresholding function, we shrink the norms according to η while keeping the phases of the row vectors. Throughout the iterations, the threshold levels γ_u and γ_v are pre-specified and remain unchanged. In order for the theorem to work, these levels can be chosen as in (9) below.

Algorithm 1: Matrix Denoising via Two-Way Iterative Thresholding

Input:

1. Observed data matrix $\mathbf{X}$.
2. Thresholding function η and thresholds γ_u and γ_v.
3. Rank r and noise standard deviation σ.
4. Initial orthonormal matrix $\mathbf{V}^{(0)} \in \mathbb{R}^{n\times r}$.

Output: Denoised matrix $\widehat{\mathbf{M}}$.

repeat

Right-to-Left Multiplication: $\mathbf{U}^{(t),\mathrm{mul}} = \mathbf{X}\mathbf{V}^{(t-1)}$.

Left Thresholding: $\mathbf{U}^{(t),\mathrm{thr}} = (u_{ij}^{(t),\mathrm{thr}})$, with $\mathbf{U}_{i*}^{(t),\mathrm{thr}} = \frac{\mathbf{U}_{i*}^{(t),\mathrm{mul}}}{\|\mathbf{U}_{i*}^{(t),\mathrm{mul}}\|}\eta(\|\mathbf{U}_{i*}^{(t),\mathrm{mul}}\|, \gamma_u)$.

Left Orthonormalization with QR Decomposition: $\mathbf{U}^{(t)}\mathbf{R}_u^{(t)} = \mathbf{U}^{(t),\mathrm{thr}}$.

Left-to-Right Multiplication: $\mathbf{V}^{(t),\mathrm{mul}} = \mathbf{X}'\mathbf{U}^{(t)}$.

Right Thresholding: $\mathbf{V}^{(t),\mathrm{thr}} = (v_{ij}^{(t),\mathrm{thr}})$, with $\mathbf{V}_{i*}^{(t),\mathrm{thr}} = \frac{\mathbf{V}_{i*}^{(t),\mathrm{mul}}}{\|\mathbf{V}_{i*}^{(t),\mathrm{mul}}\|}\eta(\|\mathbf{V}_{i*}^{(t),\mathrm{mul}}\|, \gamma_v)$.

Right Orthonormalization with QR Decomposition: $\mathbf{V}^{(t)}\mathbf{R}_v^{(t)} = \mathbf{V}^{(t),\mathrm{thr}}$.

until *Convergence*;

Compute projection matrices $\widehat{\mathbf{P}}_u = \widehat{\mathbf{U}}\widehat{\mathbf{U}}'$ and $\widehat{\mathbf{P}}_v = \widehat{\mathbf{V}}\widehat{\mathbf{V}}'$, where $\widehat{\mathbf{U}}$ and $\widehat{\mathbf{V}}$ are $\mathbf{U}^{(t)}$ and $\mathbf{V}^{(t)}$ at convergence.

Compute denoised matrix $\widehat{\mathbf{M}} = \widehat{\mathbf{P}}_u\mathbf{X}\widehat{\mathbf{P}}_v$.

To determine the convergence of the iterative part, we could either run a pre-specified number of iterations or stop after the difference between successive iterates are sufficiently small, e.g.,

$$\|\mathbf{U}^{(t)}(\mathbf{U}^{(t)})' - \mathbf{U}^{(t-1)}(\mathbf{U}^{(t-1)})'\|_{\mathrm{F}}^2 \vee \|\mathbf{V}^{(t)}(\mathbf{V}^{(t)})' - \mathbf{V}^{(t-1)}(\mathbf{V}^{(t-1)})'\|_{\mathrm{F}}^2 \le \epsilon, \tag{5}$$

where ϵ is a pre-specified tolerance level.

Initialization To initialize Algorithm 1, we need to further specify the rank r, the noise standard deviation σ and a starting point $\mathbf{V}^{(0)}$ for the iteration. For the ease of exposition, we assume that r is known. Otherwise, it can be estimated by methods such as those described in [28]. When we have Gaussian noise and $kl < \frac{1}{2}mn$, the noise standard deviation can be estimated by

$$\widehat{\sigma} = 1.4826 \cdot \mathrm{MAD}(\{\mathbf{M}_{ij} : i \in [m], j \in [n]\}). \tag{6}$$

Finally, to obtain a reasonable initial orthonormal matrix $\mathbf{V}^{(0)}$, we propose to use Algorithm 2 for the case of Gaussian noise.

Remark 1. In practice, Algorithm 1 and Algorithm 2 are not restricted to the denoising of matrices with Gaussian noise. With proper modification and robustification, they can be used together to deal with other noise distributions and/or outliers. See, e.g., [28].

3 Theoretical Results

In this section, we present a minimax theory underlying the denoising/estimation problem formulated in Section 2.1.

Algorithm 2: Initialization for Algorithm 1

Input:

1. Observed data matrix $\mathbf{X}$.
2. Tuning parameter α.
3. Rank r and noise standard deviation σ.

Output: Estimators $\widehat{\mathbf{U}} = \mathbf{U}^{(0)}$ and $\widehat{\mathbf{V}} = \mathbf{V}^{(0)}$.

1 Select the subset I_0 of rows and the subset J_0 of columns as

$$I^0 = \{i : \|\mathbf{X}_{i*}\|^2 \geq \sigma^2(n + \alpha\sqrt{n \log n})\}, \tag{7a}$$

$$J^0 = \{j : \|\mathbf{X}_{*j}\|^2 \geq \sigma^2(m + \alpha\sqrt{m \log m})\}. \tag{7b}$$

2 Compute $\mathbf{X}^{(0)} = (x_{ij}^{(0)})$, where $x_{ij}^{(0)} = x_{ij}\mathbf{1}_{i\in I^0}\mathbf{1}_{j\in J^0}$.

3 Compute $\mathbf{U}^{(0)} = [\mathbf{u}_1^{(0)}, \ldots, \mathbf{u}_r^{(0)}]$ and $\mathbf{V}^{(0)} = [\mathbf{v}_1^{(0)}, \ldots, \mathbf{v}_r^{(0)}]$, where $\mathbf{u}_\nu^{(0)}$ ($\mathbf{v}_\nu^{(0)}$) is the ν^{th} leading left (right) singular vector of $\mathbf{X}^{(0)}$.

3.1 Minimax Lower Bounds

Theorem 1. *Let $\mathcal{F} = \mathcal{F}(m, n, k, l, r, d, \kappa)$ with $\kappa > 1$ and $k \wedge l \geq 2r$. There exists a positive constant c that depends only on κ, such that for any $q \in [1, 2]$, the minimax risk for estimating $\mathbf{M}$ under the squared Schatten-q error loss (3) satisfies*

$$\inf_{\widehat{\mathbf{M}}} \sup_{\mathcal{F}} \mathbb{E} L_q(\mathbf{M}, \widehat{\mathbf{M}}) \geq c\sigma^2 \left[\left(r^{\frac{2}{q}-1}\frac{d^2}{\sigma^2}\right) \wedge \Psi_q(m, n, k, l, r)\right]$$

where the rate function $\Psi_q(m, n, k, l, r) = r^{\frac{2}{q}}(k + l) + r^{\frac{2}{q}-1}\left(k \log \frac{em}{k} + l \log \frac{en}{l}\right)$.

A proof of the theorem is given in Section 5.1.

Remark 2. Regardless of the value of q, the lower bounds reflect two different scenarios.

The first scenario is the "low signal" case where

$$d^2 \leq \sigma^2 \Psi_2(m, n, k, l, r). \tag{8}$$

In this case, the first term in the lower bound (1) dominates, and the rate is achieved by simply estimating $\mathbf{M}$ by $\mathbf{0} \in \mathbb{R}^{m\times n}$.

The second scenario is when (8) does not hold. In this case, the second term in (1) dominates. We note this term is expressed as the sum of two terms. As to be revealed by the proof, the first summand is an "oracle" error term which occurs even when the indices of the nonzero rows and columns of $\mathbf{M}$ are given by an oracle. In contrast, the second summand results from the combinatorial uncertainty about the locations of these nonzero rows and columns.

3.2 Minimax Upper Bounds

To state the upper bounds, we first specify the threshold levels used in Algorithm 1. In particular, for some sufficiently large constant $\beta > 0$, set

$$\gamma_u^2 = \gamma_v^2 = \gamma^2 = 1.01(r + 2\sqrt{r\beta \log m} + 2\beta \log m). \tag{9}$$

For Theorem 2 to hold, it suffices to choose any $\beta \geq 4$. In addition, we specify the stopping (convergence) rule for the loop in Algorithm 1. For $\mathbf{X}^{(0)}$ defined in Algorithm 2, let $d_r^{(0)}$ be its r^{th} largest singular value. Define

$$\widehat{T} = \frac{1.1}{2}\left[\frac{\log m}{\log 2} + \log \frac{(d_r^{(0)})^2}{\gamma^2}\right], \tag{10}$$

and

$$T = \frac{1.01}{2}\left[\frac{\log m}{\log 2} + \log \frac{d_r^2}{k\gamma_u^2 \vee l\gamma_v^2}\right]. \tag{11}$$

We propose to stop the iteration in Algorithm 1 after $\widehat{T}$ steps. Last but not least, we need the following technical condition.

Condition 1. There exists a sufficiently small absolute constant c, such that $m \geq n$, $\log d \leq cm$, $c \leq \log m / \log n \leq 1/c$, $\log m \leq c[(m-k) \wedge (n-l)]$, $k \vee l \leq c(m \wedge n)$. In addition, there exists a sufficiently small constant c' that depends only on κ, such that $d^{-2} r \left(k\sqrt{n \log m} + l\sqrt{m \log m}\right) \leq c'$.

With the above definition, the following theorem establishes high probability upper bounds of the proposed estimator.

Theorem 2. *Let Condition 1 be satisfied. In Algorithm 1, let $\mathbf{V}^{(0)}$ be obtained by Algorithm 2 with $\alpha \geq 4$ in (7). Let γ_u and γ_v be defined as in (9) with $\beta \geq 4$. Moreover, we stop the iteration after $\widehat{T}$ steps with $\widehat{T}$ defined in (10), and use $\widehat{\mathbf{U}} = \mathbf{U}^{(\widehat{T})}$ and $\widehat{\mathbf{V}} = \mathbf{V}^{(\widehat{T})}$ in subsequent steps. For sufficiently large values of m and n, uniformly over $\mathcal{F}(m, n, k, l, r, d, \kappa)$, with probability at least $1 - O(m^{-2})$, $\widehat{T} \in [T, 3T]$ and*

$$\|\widehat{\mathbf{M}} - \mathbf{M}\|_{\mathrm{S}_q}^2 \leq C\sigma^2 \left[r^{\frac{2}{q}}(k + l + \log m) + r^{\frac{2}{q}-1}(k+l)\log m\right]$$

where C is a positive constant that depends only on κ and β.

The proof of the theorem is given in Section 5.2.

Remark 3. Under Condition 1, for sufficient large values of m and n, (8) cannot hold, and so the relevant lower bound is $c\sigma^2\Psi_q(m, n, k, l, r)$. In comparison, when $k \wedge l \geq (1+\epsilon)r$ for any universal small constant $\epsilon > 0$, the upper bounds in Theorem 2 always matches the lower bounds for all $q \in [1, 2]$ up to a multiplicative log factor. If in addition, $\log m = O(k \vee l)$ and $k = O(m^a)$ and $l = O(n^a)$ for some constant $a \in (0, 1)$, then the rates in the lower and upper bounds match exactly for all $q \in [1, 2]$.

Remark 4. The proposed estimator is adaptive since it does not depend on the knowledge of k, l and q. Its dependence on r can also be removed, as we explain in the next subsection.

3.3 Rank Selection

We now turn to data-based selection of the rank r. Recall the sets I^0 and J^0 defined in (7). We propose to use the following data-based choice of r:

$$\widehat{r} = \max\left\{s : \sigma_s(\mathbf{X}_{I^0 J^0}) \geq \sigma\, \delta_{|I^0||J^0|}\right\}, \tag{12}$$

where for any $i \in [m]$ and $j \in [n]$, $\delta_{ij} = \sqrt{i} + \sqrt{j} + \sqrt{2i \log \frac{em}{i} + 2j \log \frac{en}{j} + 8 \log m}$. We note that it is straightforward to incorporate this rank selection step into Algorithm 2. Indeed, we can compute $\widehat{r}$ right after step 1 and replace all r in the subsequent steps by $\widehat{r}$. The following result justifies our proposal.

a	0.5	1	5	10	20
Average($L_2(\mathbf{M}, \widehat{\mathbf{M}})$)	1093.18	924.90	936.82	927.88	944.08
Standard error	(7.96)	(5.41)	(5.69)	(5.30)	(6.51)
Average($L_1(\mathbf{M}, \widehat{\mathbf{M}})$)	18346.20	15993.79	16354.86	16277.88	16526.22
Standard error	(115.06)	(84.82)	(95.22)	(89.57)	(104.87)

Table 1: Average losses (and its standard error) of $\widehat{\mathbf{M}}$ out of 100 repetitions for different choices of singular values.

Proposition 1. *Under the condition of Theorem 2, $\widehat{r} = r$ holds with probability at least $1 - O(m^{-2})$.*

A proof of the proposition is given in Section 5.3. According to Proposition 1, we can use $\widehat{r}$ as the input for rank in Algorithm 1 and the conclusion of Theorem 2 continues to hold.

4 Simulation

In this section, we demonstrate the performance of the proposed denoising method on synthetic datasets.

In the first numerical experiment, we fix $m = 2000, n = 1000, k = l = 50$ and $r = 10$. On the other hand, we set the singular values of $\mathbf{M}$ as $(d_1, \ldots, d_{10}) = a \times (200, 190, \ldots, 120, 110)$, where $a \in \{0.5, 1, 5, 10, 20\}$. The $\mathbf{U}$ matrix is obtained by orthonormalizing a $m \times r$ matrix the i^{th} row of which is filled i.i.d. $N(0, i^4)$ entries for any $i \in [k]$ and zeros otherwise. The $\mathbf{V}$ matrix is obtained in the same way with m and k replaced by n and l. The noise standard deviation is set at $\sigma = 1$. Table 1 reports the average values of $L_q(\mathbf{M}, \widehat{\mathbf{M}})$ for $q = 2$ and 1 and their standard errors out of 100 repetitions for each value of a. Throughout, we use (6) to estimate σ, Algorithm 2 with $\alpha = 4$ to compute $\mathbf{V}^{(0)}$ and (12) to select the rank. In Algorithm 1, we set $\beta = 3$ and we terminate the iteration once (5) holds with $\epsilon = 10^{-10}$. The thresholding function η is fixed to be hard thresholding $\eta(x, t) = x\mathbf{1}_{|x|>t}$. In all the repetitions, the proposed $\widehat{r}$ in (12) consistently yields the right rank $r = 10$. From the results in Table 1, we conclude that the reconstruction error is stable across different choices of the singular values of $\mathbf{M}$, which agrees well with the theoretical results in Theorem 2. We note that the magnitude of the average errors reported in Table 1 is also expected. For reference, under the simulation setting, the oracle risk term for the Schatten-2 norm, modulo a constant factor, should be $\sigma^2 r(k + l) = 1000$, and for the Schatten-1 norm, modulo a constant factor, should be $\sigma^2 r^2(k + l) = 10000$.

In the second experiment, we fix $m = 2000, n = 1000, r = 10$ and the singular values of $\mathbf{M}$ are $(d_1, \ldots, d_{10}) = (200, 190, \ldots, 120, 110)$. On the other hand, we consider four different combinations of sparsity parameters: $(k, l) = (50, 50), (50, 200), (100, 200)$ and $(100, 50)$. For each (k, l) pair, the way we generate $\mathbf{U}$, $\mathbf{V}$ and $\mathbf{X}$ is the same as that in the first experiment. Moreover, the tuning parameter values used in denoising are also the same as before. In all the repetitions, $\widehat{r}$ in (12) consistently select $r = 10$. In Table 2, we report the average values of $L_q(\mathbf{M}, \widehat{\mathbf{M}})$ for $q = 2$ and 1 and their standard errors over 100 repetitions. Moreover, we report the rescaled average loss where the rescaling constant is chosen to be $r^{\frac{2}{q}-1}(r + \log m)(k + l)$, the rate derived in Theorem 2. By the results reported in Table 2, we see that for either loss function, the rescaled average losses are stable with respect to different sparsity levels specified by different values of k and l. Again, this agrees well with the earlier theoretical results.

(k,l)	(50, 50)	(50, 200)	(100, 200)	(100, 50)
Average$(L_2(\mathbf{M},\widehat{\mathbf{M}}))$	1133.03	2662.07	3598.69	1673.49
Standard error	(5.96)	(11.73)	(12.84)	(9.73)
Average$\left(\frac{L_2(\mathbf{M},\widehat{\mathbf{M}})}{(r+\log m)(k+l)}\right)$	0.64	0.60	0.68	0.63
Average$(L_1(\mathbf{M},\widehat{\mathbf{M}}))$	19056.47	43035.95	65099.19	28347.12
Standard error	(88.42)	(172.39)	(231.98)	(146.07)
Average$\left(\frac{L_1(\mathbf{M},\widehat{\mathbf{M}})}{(r^2+r\log m)(k+l)}\right)$	1.08	0.98	1.23	1.07

Table 2: Average losses (with its standard error) and average rescaled losses of $\widehat{\mathbf{M}}$ out of 100 repetitions for different sparsity levels.

5 Proofs

5.1 Proof of Theorem 1

Proof of Theorem 1. To establish the lower bound, first consider the subset $\mathcal{F}_1 \subset \mathcal{F}(m,n,k,l,r,d,\kappa)$ where we further require $\mathrm{supp}(\mathbf{V}) = [r]$. Thus, except for the first r columns, all columns of $\mathbf{M}$ are zeros. So, by a simple sufficiency argument, we may assume that $n = l = r$. In this case, the problem of estimating $\mathbf{M}$ under model (1) can be viewed as a special case of sparse reduced rank regression where the design matrix is the identity matrix $\mathbf{I}_m$. Therefore, [19, Theorem 2] implies that

$$\inf_{\widehat{\mathbf{M}}} \sup_{\mathcal{F}} \mathbb{E} L_q(\mathbf{M},\widehat{\mathbf{M}}) \geq \inf_{\widehat{\mathbf{M}}} \sup_{\mathcal{F}_1} \mathbb{E} L_q(\mathbf{M},\widehat{\mathbf{M}}) \geq c\left[r^{\frac{2}{q}-1}d^2 \wedge \left(r^{\frac{2}{q}}k + r^{\frac{2}{q}-1}k\log\frac{em}{k}\right)\right].$$

By symmetry, we also have

$$\inf_{\widehat{\mathbf{M}}} \sup_{\mathcal{F}} \mathbb{E} L_q(\mathbf{M},\widehat{\mathbf{M}}) \geq c\left[r^{\frac{2}{q}-1}d^2 \wedge \left(r^{\frac{2}{q}}l + r^{\frac{2}{q}-1}l\log\frac{en}{l}\right)\right].$$

We complete the proof by noting that for any $a,b,c>0$, $(a\wedge b)\vee(a\wedge c) = a\wedge(b\vee c) \asymp a\wedge(b+c)$. □

5.2 Proof of Theorem 2

To prove Theorem 2, we follow the oracle sequence approach developed in [18]. Throughout the proof, we assume that $\sigma = 1$ is known. The case of general $\sigma > 0$ comes from obvious scaling arguments. In what follows, we first define the oracle sequence and introduce some preliminaries. Then we give an overview of the proof, which is divided into three steps. After the overview, the three steps are carried out in order, which then leads to the final proof of the theorem. Due to the space limit, proofs of intermediate results are omitted.

Preliminaries We first introduce some notation. For any matrix $\mathbf{A}$, $\mathrm{span}(\mathbf{A})$ stands for the subspace spanned by the column vectors of $\mathbf{A}$. If we were given the oracle knowledge of $I = \mathrm{supp}(\mathbf{U})$ and $J = \mathrm{supp}(\mathbf{V})$, then we can define an oracle version of the observed matrix as

$$\widetilde{\mathbf{X}} = (x_{ij}\mathbf{1}_{i\in I}\mathbf{1}_{j\in J}) \in \mathbb{R}^{m\times n}. \tag{13}$$

With appropriate rearrangement of rows and columns, the $I\times J$ submatrix concentrates on the top-left corner. From now on, we assume that this is the case. We denote the singular value

decomposition of $\mathbf{X}$ by

$$\widetilde{\mathbf{X}} = \begin{bmatrix} \widetilde{\mathbf{U}} & \widetilde{\mathbf{U}}_\perp \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{D}} & \mathbf{0} \\ \mathbf{0} & \widetilde{\mathbf{D}}_\perp \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{V}}' \\ (\widetilde{\mathbf{V}}_\perp)' \end{bmatrix}, \tag{14}$$

where $\widetilde{\mathbf{U}}, \widetilde{\mathbf{D}}, \widetilde{\mathbf{V}}$ consist of the first r singular triples of $\widetilde{\mathbf{X}}$, and $\widetilde{\mathbf{U}}_\perp, \widetilde{\mathbf{D}}_\perp, \widetilde{\mathbf{V}}_\perp$ contain the remaining $n-r$ triples (recall that we have assumed $m \geq n$). In particular, the successive singular values of $\widetilde{\mathbf{X}}$ are denoted by $\widetilde{d}_1 \geq \widetilde{d}_2 \geq \cdots \geq \widetilde{d}_n \geq 0$.

With the oracle knowledge of I and J, we can define oracle versions of Algorithm 2 and Algorithm 1. In the oracle version of Algorithm 2, we replace the subsets I^0 and J^0 by $\widetilde{I}^0 = I^0 \cap I$ and $\widetilde{J}^0 = J^0 \cap J$, and the output matrices are denoted by $\widetilde{\mathbf{U}}^{(0)}$ and $\widetilde{\mathbf{V}}^{(0)}$. In the oracle version of Algorithm 1, $\mathbf{X}$ is replaced by $\widetilde{\mathbf{X}}$ and $\mathbf{V}^{(0)}$ is replaced by $\widetilde{\mathbf{V}}^{(0)}$. The intermediate matrices obtained after each step within the loop are denoted by $\widetilde{\mathbf{U}}^{(t),\mathrm{mul}}$, $\widetilde{\mathbf{U}}^{(t),\mathrm{thr}}$, $\widetilde{\mathbf{U}}^{(t)}$ and $\widetilde{\mathbf{V}}^{(t),\mathrm{mul}}$, $\widetilde{\mathbf{V}}^{(t),\mathrm{thr}}$, $\widetilde{\mathbf{V}}^{(t)}$, respectively. We note that for any t, it is guaranteed that

$$\begin{aligned} \mathrm{supp}(\widetilde{\mathbf{U}}^{(t),\mathrm{thr}}) &= \mathrm{supp}(\widetilde{\mathbf{U}}^{(t)}) \subset I, \\ \mathrm{supp}(\widetilde{\mathbf{V}}^{(t),\mathrm{thr}}) &= \mathrm{supp}(\widetilde{\mathbf{V}}^{(t)}) \subset J. \end{aligned}$$

To investigate the properties of the oracle sequence, we will trace the evolution of the columns subspaces of $\widetilde{\mathbf{U}}^{(t),\mathrm{mul}}$, $\widetilde{\mathbf{U}}^{(t)}$, $\widetilde{\mathbf{V}}^{(t),\mathrm{mul}}$ and $\widetilde{\mathbf{V}}^{(t)}$. To this end, denote the r canonical angles [12] between $\mathrm{span}(\widetilde{\mathbf{U}}^{(t),\mathrm{mul}})$ and $\mathrm{span}(\widetilde{\mathbf{U}})$ by $\pi/2 \geq \phi_{u,1}^{(t)} \geq \cdots \geq \phi_{u,r}^{(t)} \geq 0$, and define

$$\sin \Phi_u^{(t)} = \mathrm{diag}(\sin \phi_{u,1}^{(t)}, \ldots, \sin \phi_{u,r}^{(t)}). \tag{15}$$

Moreover, denote the canonical angles between $\mathrm{span}(\widetilde{\mathbf{U}}^{(t)})$ and $\mathrm{span}(\widetilde{\mathbf{U}})$ by $\pi/2 \geq \theta_{u,1}^{(t)} \geq \cdots \geq \theta_{u,r}^{(t)} \geq 0$, and let

$$\sin \Theta_u^{(t)} = \mathrm{diag}(\sin \theta_{u,1}^{(t)}, \ldots, \sin \theta_{u,r}^{(t)}). \tag{16}$$

The quantities $\phi_{v,i}^{(t)}$, $\sin \Phi_v^{(t)}$, $\theta_{v,i}^{(t)}$ and $\sin \Theta_v^{(t)}$ are defined analogously. For any pair of $m \times r$ orthonormal matrices $\mathbf{W}_1$ and $\mathbf{W}_2$, let the canonical angles between $\mathrm{span}(\mathbf{W}_1)$ and $\mathrm{span}(\mathbf{W}_2)$ be $\pi/2 \geq \theta_1 \geq \cdots \geq \theta_r \geq 0$ and $\sin \Theta = \mathrm{diag}(\sin \theta_1, \ldots, \sin \theta_r)$, then [25]

$$\begin{aligned} \|\sin \Theta\|_{\mathrm{F}} &= \frac{1}{\sqrt{2}} \|\mathbf{W}_1 \mathbf{W}_1' - \mathbf{W}_2 \mathbf{W}_2'\|_{\mathrm{F}}, \\ \|\sin \Theta\|_{\mathrm{op}} &= \|\mathbf{W}_1 \mathbf{W}_1' - \mathbf{W}_2 \mathbf{W}_2'\|_{\mathrm{op}}. \end{aligned} \tag{17}$$

Overview Given the oracle sequence defined as above, we divide the proof into three steps. First, we show that the output of the oracle version of Algorithm 2 gives a good initial value for the oracle version of Algorithm 1. Next, we prove two recursive inequalities that characterize the evolution of the column subspaces of $\widetilde{\mathbf{U}}^{(t)}$ and $\widetilde{\mathbf{V}}^{(t)}$, and show that after T iterates, the output of the oracle version of Algorithm 1 estimates $\mathbf{M}$ well. Last but not least, we show that with high probability the oracle estimating sequence and the actual estimating sequence are identical up to $3T$ iterates and that $\widehat{T} \in [T, 3T]$. Therefore, the actual estimating sequence inherits all the nice properties that can be claimed for the oracle sequence.

In what follows, we carry out the three steps in order.

Initialization We first investigate the properties of $\widetilde{\mathbf{X}}$, $\widetilde{I}^0$, $\widetilde{J}^0$ and $\widetilde{\mathbf{V}}^{(0)}$.

Note that for any orthonormal matrix $\mathbf{W}$, $\mathbf{W}\mathbf{W}'$ gives the projection matrix onto span($\mathbf{W}$). The following lemma quantifies the difference between the leading singular structures of $\mathbf{X}$ and $\mathbf{M}$.

Lemma 1. *With probability at least* $1 - m^{-2}$,

$$\|\mathbf{U}\mathbf{U} - \widetilde{\mathbf{U}}\widetilde{\mathbf{U}}\|_{\mathrm{F}},\ \|\mathbf{V}\mathbf{V} - \widetilde{\mathbf{V}}\widetilde{\mathbf{V}}\|_{\mathrm{F}} \le \frac{\sqrt{2r}}{d_r}\left(\sqrt{k} + \sqrt{l} + 2\sqrt{\log m}\right), \tag{18}$$

and for any $i \in [n]$,

$$|\widetilde{d}_i - d_i| \le \sqrt{k} + \sqrt{l} + 2\sqrt{\log m} = o(d_r), \tag{19}$$

where the last equality holds under Condition 1.

Proof. By symmetry, we only need to spell out the arguments for $\mathbf{U}$ in (18). By definition, $\widetilde{\mathbf{X}} = \mathbf{U}\mathbf{D}\mathbf{V}' + \widetilde{\mathbf{Z}}$ where (after reordering of the rows and the columns) $\widetilde{\mathbf{Z}} = \begin{bmatrix} \mathbf{Z}_{IJ} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$. Thus, we have

$$\|\mathbf{U}\mathbf{U}' - \widetilde{\mathbf{U}}\widetilde{\mathbf{U}}'\|_{\mathrm{F}} \le \sqrt{2r}\|\mathbf{U}\mathbf{U}' - \widetilde{\mathbf{U}}\widetilde{\mathbf{U}}'\|_{\mathrm{op}} \le \frac{\sqrt{2r}}{d_r}\|\widetilde{\mathbf{Z}}\|_{\mathrm{op}}.$$

Here, the first inequality holds since $\mathsf{rank}(\mathbf{U}\mathbf{U}' - \widetilde{\mathbf{U}}\widetilde{\mathbf{U}}') \le 2r$ and the last inequality is due to Wedin's $\sin\theta$ theorem [27]. By the Davidson-Szarek bound [9], with probability at least $1 - m^{-2}$, $\|\widetilde{\mathbf{Z}}\|_{\mathrm{op}} = \|\mathbf{Z}_{IJ}\|_{\mathrm{op}} \le \sqrt{k} + \sqrt{l} + 2\sqrt{\log m}$. This completes the proof of (18).

On the other hand, Corollary 8.6.2 of [12] implies that $|\widetilde{d}_i - d_i| \le \|\mathbf{Z}_{IJ}\|_{\mathrm{op}}$. Together with the above discussion, we obtain the first inequality in (19). The second inequality is a direct consequence of Condition 1. This completes the proof. □

Next, we investigate the properties of the sets selected in Algorithm 2. For some universal constants $0 < a_- < 1 < a_+$, define the following two deterministic sets

$$I_{\pm}^0 = \left\{i \in [m] : \|\mathbf{M}_{i*}\|^2 \ge a_{\mp}\alpha\sqrt{n\log m}\right\}, \quad J_{\pm}^0 = \left\{j \in [n] : \|\mathbf{M}_{*j}\|^2 \ge a_{\mp}\alpha\sqrt{m\log m}\right\}. \tag{20}$$

Lemma 2. *Let Condition 1 be satisfied, and let* $\alpha \ge 4$, $a_- \le \frac{1}{20}$ *and* $a_+ \ge 2$ *be fixed constants. For sufficiently large values of* m *and* n, *with probability at least* $1 - O(m^{-2})$, *we have* $I_- \subseteq \widetilde{I}^0 \subseteq I_+$ *and* $J_- \subseteq \widetilde{J}^0 \subseteq J_+$, *and so* $I^0 = \widetilde{I}^0$ *and* $J^0 = \widetilde{J}^0$.

Proof. By symmetry, we only show the proof for $\widetilde{I}^0$ here. The arguments for $\widetilde{J}^0$ are similar. On the one hand, we have

$$\begin{aligned}
\mathbb{P}(I_-^0 \not\subseteq \widetilde{I}^0) &\le \sum_{i \in I_-^0} \mathbb{P}\left(\|\mathbf{X}_{i*}\|^2 < n + \alpha\sqrt{n\log m}\right) \\
&\le m\,\mathbb{P}\left(\chi_n^2(a_+\alpha\sqrt{n\log m}) < n + \alpha\sqrt{n\log m}\right) \\
&\le m\,\exp\left(-\frac{(a_+ - 1)^2\alpha^2 n\log m}{4n + 8a_+\alpha\sqrt{n\log m}}\right) \\
&\le m\,\exp(-3\log m) = m^{-2}.
\end{aligned}$$

Here, the last inequality holds for fixed $a_+ \ge 2$, $\alpha \ge 4$ and all sufficiently large (m, n) such that $2a_+\alpha\sqrt{n\log m} \le n/3$, which is guaranteed by Condition 1.

On the other hand, for $x = \frac{(1-a_-)^2\alpha^2 n\log m}{(2.1)^2(n+2a_-\sqrt{n\log m})}$, we have

$$\begin{aligned}
\mathbb{P}(\widetilde{I}^0 \not\subseteq I^0_+) &\le \sum_{i\in(I^0_+)^c} \mathbb{P}\left(\|\mathbf{X}_{i*}\|^2 > n+\alpha\sqrt{n\log m}\right)\\
&\le m\,\mathbb{P}\left(\chi^2_n(a_-\alpha\sqrt{n\log m}) > n+\alpha\sqrt{n\log m}\right)\\
&\le m\,\mathbb{P}\left(\chi^2_n(a_-\alpha\sqrt{n\log m}) > n+2.1\sqrt{(n+2\alpha_-\sqrt{n\log m})\,x}\right)\\
&\le m\,\mathbb{P}\left(\chi^2_n(a_-\alpha\sqrt{n\log m}) > n+2\sqrt{(n+2\alpha_-\sqrt{n\log m})x}+2x\right)\\
&\le m\,\exp(-x)\\
&\le m\exp(-3\log m) = m^{-2}.
\end{aligned}$$

Here, the fourth inequality holds for fixed $\alpha \ge 4$, $a_- \le \frac{1}{20}$, and all sufficiently large (m,n) such that $n+2a_-\sqrt{n\log m} \ge \frac{2}{0.21}(1-a_-)^2\alpha\sqrt{n\log m}$. The last inequality holds when, in addition, $0.95^2\cdot 16\cdot n \ge 3\cdot(2.1)^2\cdot(n+2a_-\sqrt{n\log m})$, which is again guaranteed by Condition 1.

Finally, when $I_- \subseteq \widetilde{I}^0 \subseteq I_+$, we have $I^0 = \widetilde{I}^0$ since $I_+ \subset I$. □

The next lemma estimates the accuracy of the starting point $\widetilde{\mathbf{V}}^{(0)}$ for the oracle version of Algorithm 1.

Lemma 3. *Let Condition 1 be satisfied, and let $\alpha \ge 4$ and $a_+ \ge 2$ be fixed constants. For sufficiently large values of m and n, uniformly over $\mathcal{F}(m,n,k,l,r,d,\kappa)$, with probability at least $1-O(m^{-2})$, for a positive constant C that depends only on κ, a_+ and α,*

$$\|\sin\widetilde{\Theta}^{(0)}_v\|_{\mathrm{F}} \le \frac{C}{d}\left[\left(r^2k^2n\log m\right)^{1/4} + \left(r^2l^2m\log m\right)^{1/4}\right] \le \frac{1}{6}.$$

Proof. Let $\mathbf{X}^{(0)}$ be the matrix defined in Step 2 of Algorithm 2, but with I^0 and J^0 replaced by $\widetilde{I}^0$ and $\widetilde{J}^0$. Then we have

$$\|\sin\widetilde{\Theta}^{(0)}_v\|_{\mathrm{F}} = \frac{1}{\sqrt{2}}\|\widetilde{\mathbf{V}}^{(0)}\widetilde{\mathbf{V}}^{(0)} - \widetilde{\mathbf{V}}\widetilde{\mathbf{V}}'\|_{\mathrm{F}} \le \frac{\sqrt{2r}}{\sqrt{2}}\|\widetilde{\mathbf{V}}^{(0)}\widetilde{\mathbf{V}}^{(0)} - \widetilde{\mathbf{V}}\widetilde{\mathbf{V}}'\|_{\mathrm{op}} \le \frac{\sqrt{r}}{\widetilde{d}_r}\|\widetilde{\mathbf{X}} - \widetilde{\mathbf{X}}^{(0)}\|_{\mathrm{op}}.$$

Here, the first equality is from (17). The second inequality holds since $\mathsf{rank}(\widetilde{\mathbf{V}}^{(0)}\widetilde{\mathbf{V}}^{(0)} - \widetilde{\mathbf{V}}\widetilde{\mathbf{V}}') \le 2r$, and the last inequality is due to Wedin's $\sin\theta$ theorem [27].

To further bound the rightmost side, we note that $\widetilde{\mathbf{X}}^{(0)}$ and $\widetilde{\mathbf{X}}$ are supported on $\widetilde{I}^0\times\widetilde{J}^0$ and $I\times J$ respectively, with $\widetilde{I}^0\times\widetilde{J}^0 \subset I\times J$. In addition, $(I\times J)\backslash(\widetilde{I}^0\times\widetilde{J}^0)$ is the union of two disjoint subsets $(I\backslash\widetilde{I}^0)\times J$ and $\widetilde{I}^0\times(J\backslash\widetilde{J}^0)$. Thus, the triangle inequality leads to

$$\begin{aligned}
\|\widetilde{\mathbf{X}} - \widetilde{\mathbf{X}}^{(0)}\|_{\mathrm{op}} &\le \|\widetilde{\mathbf{X}}_{I\backslash\widetilde{I}^0,J}\|_{\mathrm{op}} + \|\widetilde{\mathbf{X}}_{\widetilde{I}^0,J\backslash\widetilde{J}^0}\|_{\mathrm{op}}\\
&\le \|\mathbf{U}_{I\backslash\widetilde{I}^0,*}\mathbf{D}(\mathbf{V}_{J*})'\|_{\mathrm{op}} + \|\mathbf{U}_{\widetilde{I}^0*}\mathbf{D}(\mathbf{V}_{J\backslash\widetilde{J}^0*})'\|_{\mathrm{op}} + \|\mathbf{Z}_{I\backslash\widetilde{I}^0,J}\|_{\mathrm{op}} + \|\mathbf{Z}_{\widetilde{I}^0,J\backslash\widetilde{J}^0}\|_{\mathrm{op}}. \quad (21)
\end{aligned}$$

We now bound each of the four terms in (21) separately. For the first term, on the event such that the conclusion of Lemma 2 holds, we have

$$\|\mathbf{U}_{I\backslash\widetilde{I}^0,*}\mathbf{D}\mathbf{V}'_J\|_{\mathrm{op}} \le \|\mathbf{D}\|_{\mathrm{op}}\|\mathbf{V}_{J*}\|_{\mathrm{op}}\|\mathbf{U}_{I\backslash\widetilde{I}^0,*}\|_{\mathrm{op}} \le d_1\|\mathbf{U}_{I\backslash\widetilde{I}^0,*}\|_{\mathrm{F}} \le \frac{d_1}{d_r}(a_+\alpha)^{1/2}(k^2n\log m)^{1/4}.$$

Here, the last inequality is due to $I_-^0 \subset \widetilde{I}^0$, the definition of I_-^0 in (20), and the facts that $\|\mathbf{M}_{i*}\| \geq d_r\|\mathbf{U}_{i*}\|$ for all $i \in [m]$ and that $|I \backslash \widetilde{I}^0| \leq |I| \leq k$. By similar argument, on the event such that the conclusion of Lemma 2 holds, we can bound the second term in (21) as

$$\begin{aligned}\|\mathbf{U}_{\widetilde{I}^0 *}\mathbf{D}(\mathbf{V}_{J\backslash\widetilde{J}^0 *})'\|_{\mathrm{op}} &\leq \|\mathbf{U}_{\widetilde{I}^0 *}\|_{\mathrm{op}}\|\mathbf{D}\|_{\mathrm{op}}\|\mathbf{V}_{J\backslash\widetilde{J}^0 *}\|_{\mathrm{op}} \leq d_1\|\mathbf{U}\|_{\mathrm{op}}\|\mathbf{V}_{J\backslash\widetilde{J}^0 *}\|_{\mathrm{F}} \\ &\leq \frac{d_1}{d_r}(a_+\alpha)^{1/2}(l^2 m \log m)^{1/4}.\end{aligned}$$

To bound the last two terms, we first note that on the event such that the conclusion of Lemma 2 holds, both terms are upper bounded by $\|\mathbf{Z}_{IJ}\|_{\mathrm{op}}$. Together with the Davidson–Szarek bound [9], this implies that with probability at least $1 - m^{-2}$,

$$\|\mathbf{Z}_{I\backslash\widetilde{I}^0,J}\|_{\mathrm{op}} + \|\mathbf{Z}_{\widetilde{I}^0,J\backslash\widetilde{J}^0}\|_{\mathrm{op}} \leq 2\|\mathbf{Z}_{IJ}\|_{\mathrm{op}} \leq 2\left(\sqrt{k} + \sqrt{l} + 2\sqrt{\log m}\right).$$

Assembling the last five displays and observe that $\widetilde{d}_r \geq 0.9 d_r$ for sufficiently large values of (m, n) on the event such that the conclusion of Lemma 1, we obtain the first inequality in the conclusion. The second inequality is a direct consequence of Condition 1. This completes the proof. □

Evolution We now study how the column subspaces of $\widetilde{\mathbf{U}}^{(t)}$ and $\widetilde{\mathbf{V}}^{(t)}$ evolve over iterations. To this end, let

$$\rho = \widetilde{d}_{r+1}/\widetilde{d}_r, \tag{22}$$

where $\widetilde{d}_i$ denotes the i^{th} singular value of $\widetilde{\mathbf{X}}$.

Proposition 2. *For any $t \geq 1$, let $x^t = \|\sin\Theta_u^{(t)}\|_{\mathrm{F}}$, $y^t = \|\sin\Theta_v^{(t)}\|_{\mathrm{F}}$. Moreover, define*

$$\omega_u = (2\widetilde{d}_r)^{-1}\sqrt{k\gamma_u^2}, \qquad \omega_v = (2\widetilde{d}_r)^{-1}\sqrt{l\gamma_v^2}, \qquad \omega = \omega_u \vee \omega_v. \tag{23}$$

Let Condition 1 be satisfied. Then for sufficiently large values of (m, n), on the event such that the conclusions of Lemmas 1–3 hold,

1) For any $t \geq 1$, if $y^{t-1} < 1$, then

$$x^t\sqrt{1-(y^{t-1})^2} \leq \rho y^{t-1} + \omega_u, \qquad y^t\sqrt{1-(x^t)^2} \leq \rho x^t + \omega_v. \tag{24}$$

2) For any $a \in (0, 1/2]$, if

$$y^{t-1} \leq \frac{1.01\,\omega}{(1-a)(1-\rho)}, \tag{25}$$

then so is x^t. Otherwise,

$$x^t \leq y^{t-1}[1 - a(1-\rho)]. \tag{26}$$

The same conclusions hold with the ordered pair (y^{t-1}, x^t) replaced by (x^t, y^t) in (25)–(26).

Proof. 1) In what follows, we focus on showing the first inequality in (24). The second inequality follows from essentially the same argument.

Let $u^t = \|\sin\Phi_u^{(t)}\|_{\mathrm{F}}$. We first show that

$$u^t \leq \frac{\rho y^{t-1}}{\sqrt{1-(y^{t-1})^2}}. \tag{27}$$

Recall the SVD of $\widetilde{\mathbf{X}}$ in (14). In addition, let the QR factorization of $\widetilde{\mathbf{U}}^{(t),\mathrm{mul}} = \widetilde{\mathbf{Q}}^{(t)}\widetilde{\mathbf{R}}^{(t),\mathrm{mul}}$. By definition, $\widetilde{\mathbf{U}}^{(t),\mathrm{mul}} = \widetilde{\mathbf{X}}\widetilde{\mathbf{V}}^{(t-1)}$. Premultiplying both sides by $\begin{bmatrix}\widetilde{\mathbf{U}} & \widetilde{\mathbf{U}}_\perp\end{bmatrix}'$, we obtain

$$\begin{bmatrix}\widetilde{\mathbf{D}} & \mathbf{0} \\ \mathbf{0} & \widetilde{\mathbf{D}}_\perp\end{bmatrix}\begin{bmatrix}\widetilde{\mathbf{V}}'\widetilde{\mathbf{V}}^{(t-1)} \\ (\widetilde{\mathbf{V}}_\perp)'\widetilde{\mathbf{V}}^{(t-1)}\end{bmatrix} = \begin{bmatrix}\widetilde{\mathbf{U}}'\widetilde{\mathbf{Q}}^{(t)} \\ (\widetilde{\mathbf{U}}_\perp)'\widetilde{\mathbf{Q}}^{(t)}\end{bmatrix}\widetilde{\mathbf{R}}^{(t),\mathrm{mul}}.$$

In addition, let

$$\begin{bmatrix}\widetilde{\mathbf{U}}'\widetilde{\mathbf{Q}}^{(t)} \\ (\widetilde{\mathbf{U}}_\perp)'\widetilde{\mathbf{Q}}^{(t)}\end{bmatrix} = \begin{bmatrix}\mathbf{O}^{(t)} \\ \mathbf{W}^{(t)}\end{bmatrix}.$$

By the last two displays, we have

$$\mathbf{W}^{(t)} = \widetilde{\mathbf{D}}_\perp(\widetilde{\mathbf{V}}_\perp)'\widetilde{\mathbf{V}}^{(t-1)}(\widetilde{\mathbf{R}}^{(t),\mathrm{mul}})^{-1} = \widetilde{\mathbf{D}}_\perp\left[(\widetilde{\mathbf{V}}_\perp)'\widetilde{\mathbf{V}}^{(t-1)}\right]\left[\widetilde{\mathbf{V}}'\widetilde{\mathbf{V}}^{(t-1)}\right]^{-1}\widetilde{\mathbf{D}}^{-1}\left[\widetilde{\mathbf{U}}'\widetilde{\mathbf{Q}}^{(t)}\right].$$

Thus,

$$\|\mathbf{W}^{(t)}\|_{\mathrm{F}} \le \|\widetilde{\mathbf{D}}_\perp\|_{\mathrm{op}}\|(\widetilde{\mathbf{V}}_\perp)'\widetilde{\mathbf{V}}^{(t-1)}\|_{\mathrm{F}}\|[\widetilde{\mathbf{V}}'\widetilde{\mathbf{V}}^{(t-1)}]^{-1}\|_{\mathrm{op}}\|\widetilde{\mathbf{D}}^{-1}\|_{\mathrm{op}}\|\widetilde{\mathbf{U}}\|_{\mathrm{op}}\|\widetilde{\mathbf{Q}}^{(t)}\|_{\mathrm{op}}.$$

By Corollary 5.5.4 of [25], $\|\mathbf{W}^{(t)}\|_{\mathrm{F}} = u^t$, $\|(\widetilde{\mathbf{V}}_\perp)'\widetilde{\mathbf{V}}^{(t-1)}\|_{\mathrm{F}} = y^{t-1}$. Moreover, by Section 12.4.3 of [12], $\|[\widetilde{\mathbf{V}}'\widetilde{\mathbf{V}}^{(t-1)}]^{-1}\|_{\mathrm{op}} = 1/\cos\theta_{v,r}^{(t-1)} = 1/\sqrt{1-(\sin\theta_{v,r}^{(t-1)})^2} \le 1/\sqrt{1-(y^{t-1})^2}$. Here we have used the assumption that $y^{t-1} < 1$. Together with the facts that $\|\widetilde{\mathbf{D}}_\perp\|_{\mathrm{op}} = \widetilde{d}_{r+1}$, $\|\widetilde{\mathbf{D}}^{-1}\|_{\mathrm{op}} = \widetilde{d}_r^{-1}$, $\|\widetilde{\mathbf{U}}\|_{\mathrm{op}} = \|\widetilde{\mathbf{Q}}^{(t)}\|_{\mathrm{op}} = 1$, this leads to (27).

Next, we show that

$$x^t \le u^t + \frac{\omega_u}{\sqrt{1-(y^{t-1})^2}}. \tag{28}$$

To this end, let $w^t = \|\widetilde{\mathbf{Q}}^{(t)}(\widetilde{\mathbf{Q}}^{(t)})' - \widetilde{\mathbf{U}}^{(t)}(\widetilde{\mathbf{U}}^{(t)})'\|_{\mathrm{F}}$. Then, by (17) and the triangle inequality, we obtain

$$x^t \le u^t + \frac{1}{\sqrt{2}}w^t.$$

To bound w^t, note that Wedin's $\sin\theta$ theorem [27] implies

$$w^t \le \frac{\|\widetilde{\mathbf{U}}^{(t),\mathrm{mul}} - \widetilde{\mathbf{U}}^{(t)}\|_{\mathrm{F}}}{\sigma_r(\widetilde{\mathbf{U}}^{(t),\mathrm{mul}})}.$$

In the oracle version, $\widetilde{\mathbf{U}}^{(t),\mathrm{mul}}$ has at most k nonzero rows, and so $\|\widetilde{\mathbf{U}}^{(t),\mathrm{mul}} - \widetilde{\mathbf{U}}^{(t)}\|_{\mathrm{F}} \le \sqrt{k\gamma_u^2}$. For any unit vector $\mathbf{y} \in \mathrm{span}(\widetilde{\mathbf{V}}^{(t-1)})$, decompose $\mathbf{y} = \mathbf{y}_0 + \mathbf{y}_1$ where $\mathbf{y}_0 \in \mathrm{span}(\widetilde{\mathbf{V}})$ and $\mathbf{y}_1 \in \mathrm{span}(\widetilde{\mathbf{V}}_\perp)$. Then by definition, $\|\mathbf{y}_0\| \ge \cos\theta_{v,1}^{(t-1)} \ge \sqrt{1-(y^{t-1})^2}$. Thus, for any unit vector $\mathbf{x}$, $\|\widetilde{\mathbf{U}}^{(t),\mathrm{mul}}\mathbf{x}\|^2 = \|\widetilde{\mathbf{X}}\mathbf{V}^{(t-1)}\mathbf{x}\|^2 = \|\widetilde{\mathbf{X}}\mathbf{y}\|^2 = \|\widetilde{\mathbf{X}}\mathbf{y}_0\|^2 + \|\widetilde{\mathbf{X}}\mathbf{y}_1\|^2 \ge \|\widetilde{\mathbf{X}}\mathbf{y}_0\|^2 = \|\widetilde{\mathbf{X}}\widetilde{\mathbf{V}}\widetilde{\mathbf{V}}'\mathbf{y}_0\|^2 \ge (\widetilde{d}_r)^2\|\mathbf{y}_0\|^2 \ge (\widetilde{d}_r)^2[1-(y^{t-1})^2]$. Hence,

$$\sigma_r(\widetilde{\mathbf{U}}^{(t),\mathrm{mul}}) \ge \inf_{\|\mathbf{x}\|=1}\|\widetilde{\mathbf{U}}^{(t),\mathrm{mul}}\mathbf{x}\| \ge \widetilde{d}_r\sqrt{1-(y^{t-1})^2}.$$

Assembling the last three display, we obtain (28). Finally, the first inequality in (24) comes from (27), (28) and the triangle inequality.

2) Given (24), we have

$$x^t \le \frac{\rho y^{t-1} + \omega}{\sqrt{1-(y^{t-1})^2}},$$

and that $y^0 \le \frac{1}{6} \le \frac{1}{5}(1-\rho)^2$ for sufficiently large values of (m,n) due to Condition 1 and Lemma 1. The proof of part (2) then follows from the same argument as in the proof of Proposition 6.1 in [18]. □

Convergence We say that the oracle sequence has *converged* if

$$x^t \vee y^t \leq \frac{1.01\omega}{(1-m^{-1})(1-\rho)}. \tag{29}$$

This choice is motivated by the observation that $\frac{1.01\omega}{1-\rho}$ is the smallest possible value for x^t and y^t that Proposition 2 can lead to.

Proposition 3. *Let Condition 1 be satisfied and T be defined in (11). For sufficiently large values of (m,n), on the event such that the conclusions of Lemmas 1–3 hold, it takes at most T steps for the oracle sequence to converge in the sense of (29). For any t, let $\widetilde{\mathbf{P}}_u^{(t)} = \widetilde{\mathbf{U}}^{(t)}(\widetilde{\mathbf{U}}^{(t)})'$ and $\widetilde{\mathbf{P}}_v^{(t)} = \widetilde{\mathbf{V}}^{(t)}(\widetilde{\mathbf{V}}^{(t)})'$. Then there exists a constant C that depends only on κ, such that for all $t \geq T$,*

$$\|\widetilde{\mathbf{P}}_u^{(t)}\widetilde{\mathbf{X}}\widetilde{\mathbf{P}}_v^{(t)} - \widetilde{\mathbf{U}}\widetilde{\mathbf{D}}\widetilde{\mathbf{V}}\|_{\mathrm{F}}^2 \leq C\left(k\gamma_u^2 + l\gamma_v^2\right).$$

Proof. To prove the first claim, we rely on claim (2) of Proposition 2. Without loss of generality, assume that $m = 2^\nu$ for some integer $\nu \geq 1$. So $\nu = \log m/\log 2$. Let t_1 be the number of iterations needed to ensure that $x^t \vee y^t \leq \frac{1.01\omega}{(1-\frac{1}{2})(1-\rho)}$. Note that when (25) does not hold, (26) ensures that

$$y^t \leq y^{t-1}[1-a(1-\rho)]^2, \quad x^t \leq x^{t-1}[1-a(1-\rho)]^2. \tag{30}$$

Thus, it suffices to have $\left[1-\frac{1}{2}(1-\rho)\right]^{2t_1} \geq \frac{1.01\omega}{(1-\frac{1}{2})(1-\rho)}$, i.e., $2t_1|\log(1-\frac{1}{2}(1-\rho))| \geq \log(1-\frac{1}{2})(1-\rho)/(1.01\omega)$. Since $|\log(1-x)| \geq x$ for all $x \in (0,1)$, it suffices to set

$$t_1 = \frac{1}{1-\rho}\log\frac{\frac{1}{2}(1-\rho)}{1.01\omega} = \frac{1+o(1)}{2}\log\left(\frac{d_r^2}{k\gamma_u^2 \vee l\gamma_v^2}\right).$$

Next, let $t_2 - t_1$ be the number of additional iterations needed to achieve $x^t \vee y^t \leq 1.01\omega/[(1-\frac{1}{4})(1-\rho)]^2$. Before this is achieved, (30) is satisfied with $a = \frac{1}{4}$. So it suffices to have $[1-\frac{1}{4}(1-\rho)]^{2(t_2-t_1)} \leq (1-\frac{1}{2})/(1-\frac{1}{4})$, which is guaranteed if $t_2 - t_1 \geq \frac{2}{1-\rho}[\log(1-\frac{1}{4}) - \log(1-\frac{1}{2})]$. Recursively, we define t_i for $i = 3,\ldots,\nu$, such that $x^{t_i}, y^{t_i} \leq 1.01\omega/[(1-2^{-i})(1-\rho)]$. Repeating the above argument shows that it suffices to have $t_i - t_{i-1} = \frac{2^{i-1}}{1-\rho}[\log(1-2^{-i}) - \log(1-2^{-(i-1)})]$ for $i = 3,\ldots,\nu$. Therefore, if we let

$$t_\nu - t_1 = \frac{\nu+1/2}{2(1-\rho)} = \frac{(1+o(1))\log m}{2\log 2} \geq \sum_{i=1}^{\nu}\frac{2^{i-1}}{1-\rho}\left[\log(1-2^{-i}) - \log(1-2^{-(i-1)})\right],$$

then $x^t \vee y^t \leq 1.01\omega/[(1-m^{-1})(1-\rho)]$ for all $t \geq t_\nu$. We complete the proof of the first claim by noting that $T \geq t_\nu$ for sufficiently large m, n under Condition 1.

To prove the second claim, let $\widetilde{\mathbf{P}}_u = \widetilde{\mathbf{U}}\widetilde{\mathbf{U}}'$ and $\widetilde{\mathbf{P}}_v = \widetilde{\mathbf{V}}\widetilde{\mathbf{V}}'$. Then we have

$$\begin{aligned}
\|\widetilde{\mathbf{P}}_u^{(t)}\widetilde{\mathbf{X}}\widetilde{\mathbf{P}}_v^{(t)} - \widetilde{\mathbf{U}}\widetilde{\mathbf{D}}\widetilde{\mathbf{V}}\|_{\mathrm{F}} &= \|\widetilde{\mathbf{P}}_u^{(t)}\widetilde{\mathbf{X}}\widetilde{\mathbf{P}}_v^{(t)} - \widetilde{\mathbf{P}}_u\widetilde{\mathbf{X}}\widetilde{\mathbf{P}}_v\|_{\mathrm{F}} &(31)\\
&\leq \|(\widetilde{\mathbf{P}}_u^{(t)} - \widetilde{\mathbf{P}}_u)\widetilde{\mathbf{X}}\widetilde{\mathbf{P}}_v^{(t)}\|_{\mathrm{F}} + \|\widetilde{\mathbf{P}}_u\widetilde{\mathbf{X}}(\widetilde{\mathbf{P}}_v^{(t)} - \widetilde{\mathbf{P}}_v)\|_{\mathrm{F}}\\
&\leq \|\widetilde{\mathbf{P}}_u^{(t)} - \widetilde{\mathbf{P}}_u\|_{\mathrm{F}}\|\widetilde{\mathbf{X}}\|_{\mathrm{op}}\|\widetilde{\mathbf{P}}_v^{(t)}\|_{\mathrm{op}} + \|\widetilde{\mathbf{P}}_v^{(t)} - \widetilde{\mathbf{P}}_v\|_{\mathrm{F}}\|\widetilde{\mathbf{X}}\|_{\mathrm{op}}\|\widetilde{\mathbf{P}}_u\|_{\mathrm{op}}\\
&= \widetilde{d}_1\left(\|\widetilde{\mathbf{P}}_u^{(t)} - \widetilde{\mathbf{P}}_u\|_{\mathrm{F}} + \|\widetilde{\mathbf{P}}_v^{(t)} - \widetilde{\mathbf{P}}_v\|_{\mathrm{F}}\right) &(32)\\
&\leq C\left(\sqrt{k\gamma_u^2} + \sqrt{l\gamma_v^2}\right). &(33)
\end{aligned}$$

Here, the equality (31) is due to the definitions of $\widetilde{\mathbf{P}}_u$, $\widetilde{\mathbf{P}}_v$ and the fact that $\widetilde{\mathbf{U}}$, $\widetilde{\mathbf{D}}$ and $\widetilde{\mathbf{V}}$ consist of the first r singular values and vectors of $\widetilde{\mathbf{X}}$. The equality (32) holds since $\|\widetilde{\mathbf{X}}\|_{\mathrm{op}} = \widetilde{d}_1$ and $\|\widetilde{\mathbf{P}}_u\|_{\mathrm{op}} = \|\widetilde{\mathbf{P}}_v^{(t)}\|_{\mathrm{op}} = 1$ as both are projection matrices. Finally, the inequality (33) holds since $\|\widetilde{\mathbf{P}}_u^{(t)} - \widetilde{\mathbf{P}}_u\|_{\mathrm{F}} = \sqrt{2}x^t$ and $\|\widetilde{\mathbf{P}}_v^{(t)} - \widetilde{\mathbf{P}}_v\|_{\mathrm{F}} = \sqrt{2}y^t$ due to (17), the definitions in (23) and (29), and the fact that on the event such that (19) holds, $\widetilde{d}_1/\widetilde{d}_r \le 2\kappa$ when m and n are sufficiently large. This completes the proof. □

Remark 5. It is worth noting that the conclusions of Proposition 2 and Proposition 3 hold for any $\gamma_u > 0$ and $\gamma_v > 0$, though they will be used later with the specific choice of γ_u and γ_v in (9).

Proof of Upper Bounds We are now in the position to prove Theorem 2. To this end, we need to establish the equivalence between the oracle and the actual estimating sequences. The following lemma shows that with high probability, the oracle sequence and the actual sequence are identical up to $3T$ iterates.

Lemma 4. *Let γ_u and γ_v be defined as in (9) with some fixed constant $\beta \ge 4$ and let Condition 1 be satisfied. For sufficiently large m and n, with probability at least $1 - O(m^{-2})$, for all $1 \le t \le 3T$, $\mathbf{U}^{(t)}_{I^c *} = \mathbf{0}$, $\mathbf{V}^{(t)}_{J^c *} = \mathbf{0}$, and so $\mathbf{U}^{(t)} = \widetilde{\mathbf{U}}^{(t)}$ and $\mathbf{V}^{(t)} = \widetilde{\mathbf{V}}^{(t)}$.*

Proof. First of all, by Lemma 2, with probability at least $1 - O(m^{-2})$, $J^0 = \widetilde{J}^0 \subset J_+ \subset J$, and so $\widetilde{\mathbf{V}}^{(0)} = \mathbf{V}^{(0)}$. Define event $E^{(0)} = \{\mathbf{V}^{(0)} = \widetilde{\mathbf{V}}^{(0)}\}$.

We now focus on the first iteration. Define event

$$E_u^{(1)} = \left\{ \|\mathbf{Z}_{i*}\widetilde{\mathbf{V}}^{(0)}\| < \gamma_u, \forall i \in I^c \right\}.$$

On $E^{(0)} \cap E_u^{(1)}$, for any $i \in I^c$, $\mathbf{U}_{i*}^{(1),\mathrm{mul}} = \mathbf{X}_{i*}\mathbf{V}^{(0)} = \mathbf{Z}_{i*}\widetilde{\mathbf{V}}^{(0)}$. Thus, $\|\mathbf{U}_{i*}^{(1),\mathrm{mul}}\| < \gamma_u$ and so $\mathbf{U}_{i*}^{(1),\mathrm{thr}} = \mathbf{0}$ for all $i \in I^c$. This further implies $\mathbf{U}^{(1)}_{I^c *} = \mathbf{0}$ and $\mathbf{U}^{(1)} = \widetilde{\mathbf{U}}^{(1)}$. Further define event

$$E_v^{(1)} = \left\{ \|(\mathbf{Z}_{*j})'\widetilde{\mathbf{U}}^{(1)}\| < \gamma_v, \forall j \in J^c \right\}.$$

Then by similar argument, on the event $E^0 \cap E_u^{(1)} \cap E_v^{(1)}$, we have $\mathbf{V}^{(1)}_{J^c *} = \mathbf{0}$ and $\mathbf{V}^{(1)} = \widetilde{\mathbf{V}}^{(1)}$.

We now bound the probability of $(E_u^{(1)})^c$. Without loss of generality, let $J \subset [l]$. Note that for any $j \in J$, $i \in I^c$, $\widetilde{\mathbf{V}}^{(0)}$ depends on Z_{ij} only through $\|\mathbf{Z}_{I^c j}\|^2$ in the selection of $\widetilde{J}^0$ in the oracle version of Algorithm 2. Therefore, $\widetilde{\mathbf{V}}^{(0)}$ is independent of $\frac{Z_{ij}}{\|\mathbf{Z}_{I^c j}\|}$. Let $k' = |I^c|$ and $Y_1, \dots Y_l$ be i.i.d. $\chi_{k'}$ random variables independent of $\mathbf{Z}$. For any $i \in I^c$ and $j \in [l]$, let

$$\check{Z}_{ij} = Y_j \frac{Z_{ij}}{\|\mathbf{Z}_{I^c j}\|},$$

and $\check{\mathbf{Z}}_{i[l]} = (\check{Z}_{i1}, \dots, \check{Z}_{il}) \in \mathbb{R}^{1\times l}$. Since $\mathrm{supp}(\widetilde{\mathbf{V}}^{(0)}) \subset J \subset [l]$ on the event $E^{(0)}$, we obtain that for any $i \in I^c$, $\mathbf{Z}_{i*}\widetilde{\mathbf{V}}^{(0)} = \mathbf{Z}_{i[l]}\widetilde{\mathbf{V}}^{(0)} = \check{\mathbf{Z}}_{i[l]}\widetilde{\mathbf{V}}^{(0)} + (\mathbf{Z}_{i[l]} - \check{\mathbf{Z}}_{i[l]})\widetilde{\mathbf{V}}^{(0)}$. Thus,

$$\|\mathbf{Z}_{i*}\widetilde{\mathbf{V}}^{(0)}\| \le \|\check{\mathbf{Z}}_{i[l]}\widetilde{\mathbf{V}}^{(0)}_{[l]*}\| + \|(\mathbf{Z}_{i[l]} - \check{\mathbf{Z}}_{i[l]})\widetilde{\mathbf{V}}^{(0)}_{[l]*}\| \le \|\check{\mathbf{Z}}_{i[l]}\widetilde{\mathbf{V}}^{(0)}_{[l]*}\| + \|\mathbf{Z}_{i[l]} - \check{\mathbf{Z}}_{i[l]}\|\|\widetilde{\mathbf{V}}^{(0)}_{[l]*}\|_{\mathrm{op}}.$$

For the first term on the rightmost side, since $\check{\mathbf{Z}}_{i[l]}$ is independent of $\widetilde{\mathbf{V}}^{(0)}$, $\|\check{\mathbf{Z}}_{i[l]}\widetilde{\mathbf{V}}^{(0)}_{[l]*}\|^2 \sim \chi_r^2$, and so by Lemma 6, with probability at least $1 - O(m^{-\beta})$,

$$\|\check{\mathbf{Z}}_{i[l]}\widetilde{\mathbf{V}}^{(0)}_{[l]*}\|^2 \le r + 2\sqrt{\beta r \log m} + 2\beta \log m.$$

For the second term, we first note that $\|\widetilde{\mathbf{V}}^{(0)}_{[l]*}\|_{\mathrm{op}} = 1$ since it has orthonormal columns. Moreover, $\check{\mathbf{Z}}_{i[l]} - \mathbf{Z}_{i[l]} = \mathbf{Z}_{i[l]} \operatorname{diag}\left(\frac{Y_1}{\|\mathbf{Z}_{I^c 1}\|} - 1, \ldots, \frac{Y_1}{\|\mathbf{Z}_{I^c l}\|} - 1\right)$. Thus,

$$\|\check{\mathbf{Z}}_{i[l]} - \mathbf{Z}_{i[l]}\| \le \|\mathbf{Z}_{i[l]}\| \max_{j\in[l]} \left| \frac{Y_j}{\|\mathbf{Z}_{I^c j}\|} - 1 \right|.$$

By Lemma 6, with probability at least $1 - O(m^{-\beta})$,

$$\|\mathbf{Z}_{i[l]}\|^2 \le l + 2\sqrt{\beta l \log m} + 2\beta \log m.$$

By Lemma 7, for any $j \in [l]$, with probability at least $1 - O(m^{-(\beta+1)})$,

$$\left| \frac{Y_j}{\|\mathbf{Z}_{I^c j}\|} - 1 \right| \le \left| \frac{Y_j^2}{\|\mathbf{Z}_{I^c j}\|^2} - 1 \right| \le 4 \cdot 1.01 \cdot \sqrt{\frac{(\beta+1)\log m}{k'}}.$$

Here, the last inequality holds for sufficient large values of m and n, since Condition 1 implies that $(\log m)/k' = o(1)$. By the union bound, with probability at least $1 - O(m^{-\beta})$, for sufficient large values of m and n,

$$\|\check{\mathbf{Z}}_{i[l]} - \mathbf{Z}_{i[l]}\| \le 0.01\sqrt{\log m}\,,$$

since Condition 1 ensures that $l/k' = o(1)$. Assembling the last six displays, we obtain that for any $\beta > 1$, with probability at least $1 - O(m^{-\beta})$,

$$\|\mathbf{Z}_{i*}\widetilde{\mathbf{V}}^{(0)}\| \le \sqrt{r + 2\sqrt{\beta r \log m} + 2\beta \log m} + 0.01\sqrt{\log m} \le \gamma_v.$$

Applying the union bound again, we obtain that when $\beta \ge 4$ in (9),

$$\mathbb{P}\left\{(E_u^{(1)})^c\right\} \le O(m^{-3}). \tag{34}$$

Similarly, for any $j \in J^c$, $\widetilde{\mathbf{U}}^{(1)}$ depends on Z_{ij} only through $\|\mathbf{Z}_{iJ^c}\|$. Therefore, by analogous arguments, we also obtain (34) for $(E_v^1)^c$ with any fixed $\beta \ge 4$.

Turn to subsequent iterations, we further define events

$$E_u^{(t)} = \left\{\|\mathbf{Z}_{i*}\widetilde{\mathbf{V}}^{(t-1)}\| < \gamma_u,\ \forall i \in I^c\right\}, \quad E_v^{(t)} = \left\{\|(\mathbf{Z}_{*j})'\widetilde{\mathbf{U}}^{(t)}\| < \gamma_v,\ \forall j \in J^c\right\}, \quad t = 2, \ldots, 3T.$$

Iterating the above arguments, we obtain that on the event $E^{(0)} \cap (\cap_{t=1}^{3T} E_u^{(t)}) \cap (\cap_{t=1}^{3T} E_v^{(t)})$, $\mathbf{U}^{(t)}_{I^c *} = \mathbf{0}$, $\mathbf{V}^{(t)}_{J^c *} = \mathbf{0}$, and so $\mathbf{U}^{(t)} = \widetilde{\mathbf{U}}^{(t)}$ and $\mathbf{V}^{(t)} = \widetilde{\mathbf{V}}^{(t)}$. Moreover, by similar argument to that for (34), we can bound each $\mathbb{P}\{(E_u^{(t)})^c\}$ and $\mathbb{P}\{(E_v^{(t)})^c\}$ by $O(m^{-3})$ for all $t = 2, \ldots, 3T$ with any fixed $\beta \ge 4$ in (9). Finally, under Condition 1, $T = O(m)$, and so

$$\mathbb{P}\left\{E^{(0)} \cap (\cap_{t=1}^{3T} E_u^{(t)}) \cap (\cap_{t=1}^{3T} E_v^{(t)})\right\} = 1 - O(m^{-2}).$$

This completes the proof. □

Lemma 5. *Let $\widehat{T}$ be defined in (10). With probability at least $1 - O(m^{-2})$, $T \le \widehat{T} \le 3T$.*

Proof. By definition (9) and (11), we have

$$T \leq \frac{1.01}{2}\left(\frac{\log m}{\log 2} + \log \frac{d_r^2}{\gamma^2}\right).$$

On the other hand, note that $1/\log 2 \geq 1.44$ and that $\log(k \vee l) \leq \log m$ under the assumption that $m \geq n$, and hence

$$T \geq \frac{1.01}{2}\left(\frac{\log m}{\log 2} - \log m + \log \frac{d_r^2}{\gamma^2}\right) \geq \frac{1.01}{2}\left(0.44 \log m + \log \frac{d_r^2}{\gamma^2}\right).$$

On the other hand, on the event such that the conclusions of Lemmas 1–3 hold, we have

$$\begin{aligned}
|d_r^{(0)} - d_r| &\leq |d_r^{(0)} - \widetilde{d}_r| + |\widetilde{d}_r - d_r| \\
&= |\widetilde{d}_r^{(0)} - \widetilde{d}_r| + |\widetilde{d}_r - d_r| \\
&\leq \|\widetilde{\mathbf{X}}^{(0)} - \widetilde{\mathbf{X}}\|_{\mathrm{op}} + o(d_r) \\
&= o(d_r).
\end{aligned}$$

Hence for sufficiently large values of m and n, $\log \gamma^2 > 1$ and with probability at least $1 - O(m^{-2})$, $|\log(d_r^{(0)})^2 / \log d_r^2 - 1| \leq 0.01$. When the above inequalities all hold, we obtain $\widehat{T} \in [T, 3T]$. □

We are now in the position to prove Theorem 2.

Proof of Theorem 2. Note that on the events such that the conclusions of Lemmas 1–5 hold, we have

$$\begin{aligned}
&\|\widehat{\mathbf{M}} - \mathbf{M}\|_{\mathrm{F}} \\
&\quad = \|\widehat{\mathbf{P}}_u^{(\widehat{T})} \mathbf{X} \widehat{\mathbf{P}}_v^{(\widehat{T})} - \mathbf{U}\mathbf{D}\mathbf{V}'\|_{\mathrm{F}} \\
&\quad \leq \|\widehat{\mathbf{P}}_u^{(\widehat{T})} \mathbf{X} \widehat{\mathbf{P}}_v^{(\widehat{T})} - \widetilde{\mathbf{U}}\widetilde{\mathbf{D}}\widetilde{\mathbf{V}}'\|_{\mathrm{F}} + \|\widetilde{\mathbf{U}}\widetilde{\mathbf{D}}\widetilde{\mathbf{V}}' - \mathbf{U}\mathbf{D}\mathbf{V}'\|_{\mathrm{F}} \\
&\quad = \|\widetilde{\mathbf{P}}_u^{(\widehat{T})} \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v^{(\widehat{T})} - \widetilde{\mathbf{P}}_u \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v\|_{\mathrm{F}} + \|\widetilde{\mathbf{P}}_u \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v - \mathbf{P}_u \mathbf{M} \mathbf{P}_v\|_{\mathrm{F}} \\
&\quad \leq \|\widetilde{\mathbf{P}}_u^{(\widehat{T})} \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v^{(\widehat{T})} - \widetilde{\mathbf{P}}_u \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v\|_{\mathrm{F}} + \|\widetilde{\mathbf{P}}_u \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v - \widetilde{\mathbf{P}}_u \mathbf{M} \widetilde{\mathbf{P}}_v\|_{\mathrm{F}} \\
&\qquad + \|\widetilde{\mathbf{P}}_u \mathbf{M} \widetilde{\mathbf{P}}_v - \mathbf{P}_u \mathbf{M} \mathbf{P}_v\|_{\mathrm{F}}.
\end{aligned}$$

Here, the first and the second inequalities are both due to the triangle inequality. The second equality is due to Lemma 5 and the facts that $\mathrm{supp}(\widetilde{\mathbf{U}}^{(t)}) \subset I$, $\mathrm{supp}(\widetilde{\mathbf{V}}^{(t)}) \subset J$ and that $\widetilde{\mathbf{U}}$ and $\widetilde{\mathbf{V}}$ collect the first r left and right singular vectors of $\widetilde{\mathbf{X}}$.

We now bound each of the three terms on the rightmost side of the last display. First, on the event such that the conclusions of Proposition 3 and Lemma 5 hold, we have

$$\|\widetilde{\mathbf{P}}_u^{(\widehat{T})} \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v^{(\widehat{T})} - \widetilde{\mathbf{P}}_u \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v\|_{\mathrm{F}} \leq C\sqrt{k\gamma_u^2 + l\gamma_v^2}.$$

Next, by similar argument to that leading to the conclusion of Lemma 3, with probability at least $1 - O(m^{-2})$

$$\begin{aligned}
\|\widetilde{\mathbf{P}}_u \widetilde{\mathbf{X}} \widetilde{\mathbf{P}}_v - \widetilde{\mathbf{P}}_u \mathbf{M} \widetilde{\mathbf{P}}_v\|_{\mathrm{F}} &\leq \|\widetilde{\mathbf{X}} - \mathbf{M}\|_{\mathrm{F}} = \|\mathbf{Z}_{IJ}\|_{\mathrm{F}} \\
&\leq \sqrt{r}\|\mathbf{Z}_{IJ}\|_{\mathrm{op}} \\
&\leq \sqrt{r}(\sqrt{k} + \sqrt{l} + 2\sqrt{\log m}).
\end{aligned}$$

Last but not least,

$$\begin{aligned}
&\|\widetilde{\mathbf{P}}_u\mathbf{M}\widetilde{\mathbf{P}}_v - \mathbf{P}_u\mathbf{M}\mathbf{P}_v\|_{\mathrm{F}} \\
&\quad\le \|(\widetilde{\mathbf{P}}_u - \mathbf{P}_u)\mathbf{M}\widetilde{\mathbf{P}}_v\|_{\mathrm{F}} + \|\mathbf{P}_u\mathbf{M}(\widetilde{\mathbf{P}}_v - \mathbf{P}_v)\|_{\mathrm{F}} \\
&\quad\le d_1(\|\widetilde{\mathbf{P}}_u - \mathbf{P}_u\|_{\mathrm{F}} + \|\widetilde{\mathbf{P}}_v - \mathbf{P}_v\|_{\mathrm{F}}) \\
&\quad\le \kappa\sqrt{r}(\sqrt{k} + \sqrt{l} + 2\sqrt{\log m}).
\end{aligned}$$

Assembling the last four displays, we complete the proof for the case of Frobenius norm, i.e., $q = 2$. To obtain the result for all $q \in [1, 2)$, simply note that for any matrix $\mathbf{A}$, $\|\mathbf{A}\|_{\mathrm{s}_q} \le (\mathsf{rank}(\mathbf{A}))^{\frac{1}{q}-\frac{1}{2}}\|\mathbf{A}\|_{\mathrm{F}}$ and that $\mathsf{rank}(\widetilde{\mathbf{M}} - \mathbf{M}) \le 2r$. This completes the proof. □

5.3 Proof of Proposition 1

Proof of Proposition 1. Without loss of generality, assume that $\sigma = 1$. We first show that $\widehat{r} \le r$ with probability at least $1 - O(m^{-2})$. To this end, note that

$$\begin{aligned}
\mathbb{P}\{\widehat{r} > r\} &= \mathbb{P}\left\{\sigma_{r+1}(\mathbf{X}_{I^0J^0}) > \delta_{|I^0||J^0|}\right\} \\
&\le \mathbb{P}\left\{\max_{|A|=|I^0|,|B|=|J^0|} \sigma_{r+1}(\mathbf{X}_{AB}) > \delta_{|A||B|}\right\} \\
&\le \sum_{i=r+1}^{m}\sum_{j=r+1}^{n} \mathbb{P}\left\{\max_{|A|=i,|B|=j} \sigma_{r+1}(\mathbf{X}_{AB}) > \delta_{ij}\right\}.
\end{aligned}$$

By the interlacing property of singular values, we know that for $\mathbf{Z}$, a $m \times n$ standard Gaussian random matrix,

$$\max_{|A|=i,|B|=j} \sigma_{r+1}(\mathbf{X}_{AB}) \overset{st}{<} \max_{|A|=i-r,|B|=j-r} \sigma_1(\mathbf{Z}_{AB}) \overset{st}{<} \max_{|A|=i,|B|=j} \sigma_1(\mathbf{Z}_{AB}),$$

where $\overset{st}{<}$ means stochastically smaller. Together with the union bound, this implies

$$\begin{aligned}
\mathbb{P}\left\{\max_{|A|=i,|B|=j} \sigma_{r+1}(\mathbf{X}_{AB}) > \delta_{ij}\right\} &\le \binom{m}{i}\binom{n}{j}\mathbb{P}\{\sigma_1(\mathbf{Z}_{AB}) > \delta_{ij}\} \\
&\le \left(\frac{\mathrm{e}m}{i}\right)^i \left(\frac{\mathrm{e}n}{j}\right)^j \exp\left(-i\log\frac{\mathrm{e}m}{i} - j\log\frac{\mathrm{e}m}{j} - 4\log m\right) \\
&= m^{-4}.
\end{aligned}$$

Here, the second inequality is due to $\binom{p}{k} \le (\mathrm{e}p/k)^k$ for any $k \in [p]$ and the Davidson-Szarek bound [9]. As $n \le m$ under Condition 1, we obtain

$$\mathbb{P}\{\widehat{r} > r\} \le \sum_{i=r+1}^{m}\sum_{j=r+1}^{n} m^{-4} \le m^{-2}.$$

To show that $\widehat{r} \ge r$ with probability at least $1 - O(m^{-2})$, we note that on the event such that the conclusions of Lemmas 1–3 hold, $\sigma_r(\mathbf{X}_{I^0J^0}) = \sigma_r(\widetilde{\mathbf{X}}^0) = \widetilde{d}_r^{(0)}$. So by the triangle inequality, the conclusion of Lemma 1 and the proof of Lemma 3, we obtain that

$$\sigma_r(\mathbf{X}_{I^0J^0}) = \widetilde{d}_r^{(0)} \ge d_r - \|\mathbf{X} - \widetilde{\mathbf{X}}\|_{\mathrm{op}} - \|\widetilde{\mathbf{X}} - \widetilde{\mathbf{X}}^0\|_{\mathrm{op}} \ge d_r/4 > \delta_{kl},$$

where the second last and the last inequalities hold under Condition 1 for sufficiently large values of m and n. Note that on the event such that the conclusion of Lemma 2 holds, we have $|I^0| \le k$ and $|J^0| \le l$ and so $\delta_{|I^0||J^0|} \le \delta_{kl}$. This completes the proof. □

A Appendix

Lemma 6 (Lemma 8.1 in [1]). *Let X follow the non-central chi square distribution $\chi^2_\nu(\delta)$ with degrees of freedom ν and non-centrality parameter $\delta \geq 0$. Then for any $x > 0$,*

$$\mathbb{P}\left\{X \geq \nu + \delta + 2\sqrt{(\nu + 2\delta)x} + 2x\right\} \leq \mathrm{e}^{-x},$$
$$\mathbb{P}\left\{X \leq \nu + \delta - 2\sqrt{(\nu + 2\delta)x}\right\} \leq \mathrm{e}^{-x}.$$

Lemma 7. *Let X and Y be two independent χ^2_ν random variables. Then for any $x > 0$,*

$$\mathbb{P}\left\{\left|\frac{X}{Y} - 1\right| \leq \frac{4\sqrt{x/\nu}(1 + \sqrt{x/\nu})}{1 - 2\sqrt{x/\nu}}\right\} \geq 1 - 4\mathrm{e}^{-x}.$$

Proof. By the triangle inequality,

$$\left|\frac{X}{Y} - 1\right| \leq \frac{1}{|Y|}\left(|X - \nu| + |Y - \nu|\right).$$

By Lemma 6, for any $x > 0$, each of the following holds with probability at least $1 - 2\mathrm{e}^{-x}$:

$$|X - \nu| \leq 2\sqrt{\nu x} + 2x,$$
$$|Y - \nu| \leq 2\sqrt{\nu x} + 2x, \quad \text{and} \quad |Y| \geq \nu - 2\sqrt{\nu x}.$$

Assembling the last two displays, we complete the proof. □

References

[1] L. Birgé. *An alternative point of view on Lepski's method*, volume 36 of *Lecture Notes-Monograph Series*, pages 113–133. Institute of Mathematical Statistics, 2001.

[2] A. Buja, Z. Ma, and D. Yang. Optimal denoising of simultaneously sparse and low rank matrices in high dimensions. In *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*, pages 445–447. IEEE, 2013.

[3] F. Bunea, Y. She, and M.H. Wegkamp. Optimal selection of reduced rank estimators of high-dimensional matrices. *The Annals of Statistics*, 39(2):1282–1309, 2011.

[4] F. Bunea, Y. She, and M.H. Wegkamp. Joint variable and rank selection for parsimonious estimation of high dimensional matrices. *The Annals of Statistics*, 40:2359–2388, 2012.

[5] C. Butucea and Yu.I. Ingster. Detection of a sparse submatrix of a high-dimensional noisy matrix. *arXiv preprint arXiv:1109.0898*, 2011.

[6] E.J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6): 925–936, 2010.

[7] E.J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.

[8] G. Chen, P.F. Sullivan, and M.R. Kosorok. Biclustering with heterogeneous variance. *Proceedings of the National Academy of Sciences*, 110(30):12253–12258, 2013.

[9] K.R. Davidson and S. Szarek. *Handbook on the Geometry of Banach Spaces*, volume 1, chapter Local operator theory, random matrices and Banach spaces, pages 317–366. Elsevier Science, 2001.

[10] D.L. Donoho and M. Gavish. Minimax risk of matrix denoising by singular value thresholding. *arXiv preprint arXiv:1304.2085*, 2013.

[11] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348–1360, 2001.

[12] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins University Press, 3rd edition, 1996.

[13] R.H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *The Journal of Machine Learning Research*, 11:2057–2078, 2010.

[14] V. Koltchinskii, K. Lounici, and A.B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.

[15] M. Lee, H. Shen, J.Z. Huang, and J.S. Marron. Biclustering via sparse singular value decomposition. *Biometrics*, 66:1087–1095, 2010.

[16] X. Li and V. Voroninski. Sparse signal recovery from quadratic measurements via convex programming. *arXiv preprint arXiv:1209.4785*, 2012.

[17] K. Lounici, M. Pontil, S. Van De Geer, and A. B. Tsybakov. Oracle inequalities and optimal inference under group sparsity. *The Annals of Statistics*, 39(4):2164–2204, 2011.

[18] Z. Ma. Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*, 41(2):772–801, 2013.

[19] Z. Ma and T. Sun. Adaptive sparse reduced-rank regression. *arXiv preprint arXiv:1403.1922*, 2014.

[20] S. Negahban and M.J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, 39(2):1069–1097, 2011.

[21] S. Oymak and B. Hassibi. Asymptotically exact denoising in relation to compressed sensing. *arXiv preprint arXiv:1305.2714*, 2013.

[22] S. Oymak, A. Jalali, M. Fazel, Y.C. Eldar, and B. Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *arXiv preprint arXiv:1212.3753*, 2012.

[23] A.A. Shabalin, V.J. Weigman, C.M. Perou, and A.B. Nobel. Finding large average submatrices in high dimensional data. *The Annals of Applied Statistics*, 3:985–1012, 2009.

[24] Y. Shechtman, Y.C. Eldar, A. Szameit, and M. Segev. Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing. *Optics Express*, 19(16): 14807–14822, 2011.

[25] G.W. Stewart and J.-G. Sun. *Matrix Perturbation Theory*. Computer science and scientific computing. Academic Press, 1990.

[26] X. Sun and A.B. Nobel. On the maximal size of large-average and anova-fit submatrices in a gaussian random matrix. *Bernoulli*, 19(1):275–294, 2013.

[27] P.-A. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT*, 12: 99–111, 1972.

[28] D. Yang, Z. Ma, and A. Buja. A sparse SVD method for high-dimensional data. *arxiv preprint arXiv:1112.2433*, 2011.

[29] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B*, 68(1):49–67, 2006.

[30] X.-T. Yuan and T. Zhang. Truncated power method for sparse eigenvalue problems. *Journal of Machine Learning Research*, 14:899–925, 2013.

[31] C.H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.